

Europäisches Patentamt
European Patent Office
Office européen des brevets



(11) EP 1 215 624 A2

(12) EUROPEAN PATENT APPLICATION

(43) Date of publication:
19.06.2002 Bulletin 2002/25

(51) Int Cl.7: G06T 1/00

(21) Application number: 01204593.6

(22) Date of filing: 30.11.2001

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE TR
Designated Extension States:
AL LT LV MK RO SI

(72) Inventors:
• Jones, Paul William
Rochester, New York 14650-2201 (US)
• Honsinger, Chris W.
Rochester, New York 14650-2201 (US)

(30) Priority: 13.12.2000 US 736067


(74) Representative:
Nunney, Ronald Frederick Adolphe et al
Kodak Limited,
Patent Department (W92)-3A,
Headstone Drive
Harrow, Middlesex HA1 4TY (GB)

(71) Applicant: EASTMAN KODAK COMPANY
Rochester, New York 14650-2201 (US)

(54) System and method for embedding a watermark signal that contains message data in a digital image

(57) A method for embedding a watermark signal that contains message data in a digital image represented as a two-dimensional array of pixel values, includes the steps of: providing a dispersed message image having pixel values representative of the message data;

modifying each pixel value of the dispersed message image as a function of the corresponding pixel value in the digital image; and combining the modified dispersed message image with the digital image to produce a watermarked image.



EP 1 215 624 A2

Description

[0001] The invention relates generally to the field of digital image processing, and in particular to a method for embedding watermarks in digital images.

5 [0002] Digital watermarking refers to the embedding of a message in an image or image sequence for such purposes as establishing ownership, tracking the origin of the data, preventing unauthorized copying, or conveying additional information (meta-data) about the content. Watermarking has potential uses in a wide range of products, including digital still and video cameras, printers and other hardcopy output devices, and content delivery services (e.g., Internet-based photofinishing).

10 [0003] Numerous image watermarking methods have been described in the prior art, including both patents and the technical literature. Many of these methods are described in review papers such as: Hartung and Kutter, *Multimedia Watermarking Techniques*, Proc. IEEE, 87(7), pp. 1079-1107 (1999), and Wolfgang et al., *Perceptual Watermarks for Digital Images and Video*, "Proc. IEEE, 87(7), pp. 1108-1126 (1999).

15 [0004] A basic distinction between various methods is whether the watermark is applied in the spatial domain or the frequency domain. In either approach, many techniques use a pseudo-random (PN) sequence in the watermark generation and extraction processes. The PN sequence serves as a carrier signal, which is modulated by the original message data, resulting in dispersed message data (i.e., the watermark) that is distributed across a number of pixels in the image. A secret key (i.e., seed value) is commonly used in generating the PN sequence, and knowledge of the key is required to extract the watermark and the associated original message data.

20 [0005] The use of a PN sequence in the watermarking process typically results in a watermark signal that has similar visual characteristics to a random noise field. While noise is an inherent component of digital images because of sensor noise (whether originating from film grain when film is scanned or from sensor noise in a digital camera), it is generally desirable to minimize the additional noise that is introduced by the watermarking process. This is accomplished by reducing the amplitude of the watermark signal so that it is below the threshold of perception (i.e. subthreshold). The result is a watermark that is visually transparent, i.e. the watermark cannot be perceived by an observer under typical viewing conditions.

25 [0006] In some applications it may be impossible to provide sufficient robustness to the embedded watermark and still meet the goal of visual transparency. This is because a low amplitude watermark is more vulnerable to removal attacks, where an attempt is made to remove the watermark by using various types of image processing methods, e.g. lowpass filtering. In such cases, it becomes necessary to increase the watermark amplitude, which may result in a watermark pattern that is easily detectable (i.e. suprathreshold), and perhaps even objectionable, when viewed by an observer.

30 [0007] There are watermarking techniques that intentionally use suprathreshold watermarks. Some of these techniques embed watermarks that are easily visible to any observer, e.g., a logo that is placed in the corner of a television screen to indicate the broadcasting station. These types of watermarks are called obvious watermarks. Other suprathreshold watermark techniques are more subtle in that the watermark may be obvious to an observer who knows the nature of the watermark, but it is not obvious to an uninformed observer. An example of this latter approach can be found in U.S. Patent No. 5,699,427 issued December 16, 1997 to Chow et al. entitled *Method to Deter Document and Intellectual Property Piracy Through Individualization*. This patent teaches a method for modifying documents and images in such a way that numerous variations of the original can be produced, but these variations do not change the overall meaning or usefulness of the documents or images. The variations can include changing the spelling of words in a document or the color of objects in an image. A limitation of this approach is that the amount of embedded information is constrained, and it is also very difficult to embed arbitrary information such as a time/date stamp. Moreover, the method requires significant manual intervention to identify the words, objects, etc. that can be modified without affecting the overall meaning or usefulness of the document or image.

45 [0008] There is a need therefore to have a watermarking technique for images that allows the embedding of a high amplitude watermark in such a way that: 1) the resulting watermark is not objectionable to an observer; 2) arbitrary information can be included in the watermark; and 3) manual intervention is not required to perform the embedding process.

50 [0009] The need is met according to the present invention by providing a method for embedding a watermark signal that contains message data in a digital image represented as a two-dimensional array of pixel values that includes the steps of: providing a dispersed message image having pixel values representative of the message data; modifying each pixel value of the dispersed message image as a function of the corresponding pixel value in the digital image; and combining the modified dispersed message image with the digital image to produce a watermarked image.

55 [0010] The present invention has the advantage of making the visual characteristics of the watermark look like an expected attribute of an image that is produced by a given imaging system. In particular, the present invention minimizes the objectionability of the watermark pattern by making its appearance similar to the expected noise for the given imaging system.

Fig. 1 is a prior art method for embedding a watermark in an original image;
 Fig. 2 is a prior art method for extracting a watermark from an image containing an embedded watermark;
 Fig. 3 is a block diagram that illustrates the embedding of a signal-dependent watermark according to the present invention;
 Fig. 4 is an example of noise standard deviation as a function of signal level for signal-dependent noise;
 Fig. 5 is a prior art preferred embodiment for the frequency spectrum of a watermark carrier signal;
 Fig. 6 is an example of the frequency spectrum of film grain noise;
 Fig. 7 is an example of a carrier frequency spectrum that is modified to simulate the frequency spectrum of film grain noise; and
 Fig. 8 is an example of the frequency spectrum components that represent the carrier support region and the film grain noise support region.

[0011] The present invention is preferably implemented by a programmed digital computer. The computer can be a general purpose digital computer or a special purpose computer for digital image processing. It is within the ordinary skill in the programming art to provide a computer program for practicing the present invention from the following description of the invention.

[0012] A preferred data embedding technique for use with the present invention is disclosed in U.S. Patent No. 6,044,156 issued March 28, 2000 to Honsinger et al. entitled *Method for Generating an Improved Carrier for Use in an Image Data Embedding Application*. This patent is included in its entirety by reference. Referring to Fig. 1, in this technique, an original two-dimensional digital image 10, $I(x,y)$, is processed to produce a watermarked image 12, $I'(x,y)$. A two-dimensional message 14, $M(x,y)$, represents the data to be embedded in the original image 10. In its most general form, the message 14 is an image, and it can represent an icon 16 (e.g., a trademark), or it can represent the bits 18 in a binary message. In the latter case, the on and off states of the bits are represented as plus and minus ones (more specifically, positive and negative delta functions), which are placed in predefined and unique locations across the message image. Examples of iconic message data are trademarks, corporate logos or other arbitrary images. In order to minimize the message energy, an edge map of the icon is often used instead of the actual icon. Examples of binary message data are 32-bit representations of URL's, and copyright ID codes, or authentication information.

[0013] As shown in Fig. 1, the fundamental steps for embedding message data in an original image with this method are:

1. A $n \times n$ message image 14, $M(x,y)$, is generated from the message data;
2. The message image 14 is circularly convolved 20 with a $n \times n$ carrier image 22, $C(x,y)$ to produce a $n \times n$ dispersed message image 24. The carrier image may be produced using a secure key 26 as is known in the prior art;
3. The dispersed message image 24 is scaled 28 in amplitude using a multiplicative factor α ; and
4. The scaled dispersed message image 30 is added to the original image 10 as contiguous $n \times n$ tiles to form a watermarked image 12, $I'(x,y)$.

[0014] The carrier image 22 is derived from a PN sequence, and its properties can be modified to achieve certain performance goals, such as robustness to removal attacks. The scaling factor α is an arbitrary constant that is generally chosen to make the embedded energy simultaneously invisible and robust to common processing. Typically, the size of the dispersed message image 24 is chosen to be smaller than the size of original image 10, and the tiling process allows the scaled dispersed message image 30 to be repetitively embedded over the extent of the original image 10. The repetitive structure provides robustness to the watermark when image processing operations (such as cropping, compression, lowpass filtering, etc.) are applied to the watermarked image. Other watermarking techniques use different methods for embedding the message data, but the repetitive nature of the embedding process is a common aspect because of this improved robustness.

[0015] This embedding process for each tile can be described mathematically as:

$$I'_i(x,y) = \alpha[M(x,y)*C(x,y)] + I_i(x,y), \quad (1)$$

where the symbol $*$ represents circular convolution and the subscript i indicates the i^{th} tile ($i = 1, \dots, N$). From Fourier theory, spatial convolution is equivalent in the frequency domain to adding phase while multiplying magnitudes. Therefore, the effect of convolving the message image 14 with the carrier image 22 is to distribute the message energy in accordance with the phase of the carrier image and to modulate the amplitude spectrum of the message image with the amplitude spectrum of the carrier image. If the message image were a single delta function, $\delta(x,y)$, and the carrier image had random phase and uniform Fourier magnitude, the effect of convolving with the carrier image would be to

distribute the delta function over space. Similarly, the effect of convolving a message image with a random phase carrier image is to spatially disperse the message energy.

[0016] As shown in Fig. 2, the process as described by Honsinger et al. for extracting the message data from a watermarked image 12 consists of the following fundamental steps:

1. Contiguous $n \times n$ tiles 12' are formed from the watermarked image 12, $I'(x,y)$;
2. The tiles 12' are averaged 32 across each spatial location (x,y) to form an averaged tile 34;
3. The averaged tile 34 is circularly correlated 36 with the $n \times n$ carrier image 22 $C(x,y)$ to produce an extracted $n \times n$ message image 14', $M(x,y)$; and
4. The message data is recovered from the extracted message image 14'.

[0017] The averaging 32 of the individual tiles 12' improves the extraction of the message image because the dispersed message image in each tile will add constructively (since it is the same in each tile), while the corresponding original image content in each tile will add destructively (since it is typically different in each tile).

[0018] This watermark extraction process can be described mathematically as:

$$\begin{aligned}
 M'(x,y) &= \left[\frac{1}{N} \sum_{i=1}^N I'_i(x,y) \right] \otimes C(x,y) \\
 &= \left[\frac{1}{N} \sum_{i=1}^N (\alpha[M(x,y) * C(x,y)] + I_i(x,y)) \right] \otimes C(x,y) \\
 &= \alpha[M(x,y) * C(x,y)] \otimes C(x,y) + \left[\frac{1}{N} \sum_{i=1}^N I_i(x,y) \right] \otimes C(x,y)
 \end{aligned} \tag{2}$$

where the symbol \otimes represents circular correlation. Correlation is similar to convolution in that Fourier magnitudes also multiply. In correlation, however, phase subtracts. Therefore, the phase of the carrier image subtracts when the watermarked image is correlated with the carrier image, thus leaving the message image. Indeed, if we assume that the carrier image is designed to have uniform Fourier amplitude, then the process of correlation of the carrier image on the watermarked image in Eq. 2, can be reduced to:

$$M'(x,y) = \alpha M(x,y) + \text{noise}. \tag{3}$$

[0019] That is, the extracted message image is a scaled version of the original message image plus noise due to the cross correlation of the summed original image tiles with the carrier image.

[0020] In the method of Honsinger et al., the amplitude of the watermark is adjusted by changing the scaling factor α . However, the amplitude is always adjusted by the same amount throughout the entire image. Referring to Fig. 3, in the present invention, the amplitude of the dispersed message image 24 is scaled 28' at each spatial location (x,y) in each tile i to produce a scaled dispersed message image 30'. With this modification, the embedding equation becomes:

$$I'_i(x,y) = \alpha_i(x,y)[M(x,y) * C(x,y)] + I_i(x,y), \tag{4}$$

where $\alpha_i(x,y)$ is a signal-dependent scaling factor that varies from pixel to pixel within a given tile and from tile to tile. Specifically,

$$\alpha_i(x,y) = f(I_i(x,y)), \tag{5}$$

where $f(I_i(x,y))$ indicates that the dispersed message image amplitude scaling factor is a function of the original image pixel value I_i at each spatial location (x,y) .

[0021] As an illustration of signal-dependent noise and the resulting effect on the watermark amplitude scaling factor,

we consider a system that uses a conventional color negative photographic film as the original capture medium. After exposure and development, the film is digitized by scanning with a digital scanner. Both film grain and scanner noise are introduced during this process, and we can determine the overall system noise as a function of the signal level (e.g., the signal level could be the codevalue that corresponds to a particular density in the original negative). The characterization of the system noise is typically performed by measuring the standard deviation, σ , of an image region containing a uniform signal level, and repeating this measurement at different signal levels. A complete description of methods for characterizing noise in film-based systems can be found in U.S. Patent No. 5,641,596 issued June 24, 1997 to Gray et al. entitled *Adjusting Film Grain Properties in Digital Images*.

[0022] An example of the noise standard deviation as a function of signal level is shown in Fig. 4, where all values have been specified in terms of 8-bit codevalues. In this figure, three different single-dependent relationships are shown: one for the red color channel 40, one for the green channel 42, and one for the blue channel 44. Regardless of the color channel, it can be seen the noise standard deviation first increases slightly, then decreases, and finally increases again as the signal level is increased. This behavior is the result of the combination of the film grain noise increasing and then decreasing as the signal level increases, while the scanner noise starts low and then increases with increasing signal level.

[0023] To apply the signal-dependent characterization of the system noise as shown in Fig. 4, the scaling factor is made proportional to the noise standard deviation that corresponds to the signal level at each location, i.e.

$$\alpha_i(x,y) = k\sigma(l_i(x,y)), \quad (6)$$

where k is a constant. We can choose k so that the amplitude of the resulting watermark signal matches the overall system noise, or it may advantageous to increase or decrease k to provide more or less noise, respectively, in the watermarked image. The mapping of the image values $l_i(x, y)$ to a noise standard deviation σ can be performed using a lookup table or a mathematical equation. Such mappings are described in the aforementioned patent by Gray et al. To better mimic the overall system noise characteristics, it is also desirable to use separate mappings for the red, green, and blue color channels as shown in Fig. 4. For this example, the result is that the watermarked image will have more noise in the blue channel than the red and green channels. This is a common characteristic of film/scanner systems.

[0024] There are watermarking techniques in the prior art that also adjust the watermark amplitude based upon local signal content in the original image. Descriptions of such techniques can be found, for example, in the aforementioned review paper by Wolfgang et al.; in a paper by Kundar et al. entitled *A Robust Digital Image Watermarking Method Using Wavelet-Based Fusion*, Proc. IEEE Int. Conf. Image Processing (ICIP), Vol. 1, pp. 544-547, 1997; and in U.S. Patent No. 6,031,914 issued February 29, 2000 to Tewfik et al. entitled *Method and Apparatus for Embedding Data, Including Watermarks, in Human Perceptible Images*. There are two important distinctions between these prior art methods and the present invention. First, the motivation for locally changing the watermark amplitude in the prior art methods is to take advantage of perceptual masking. Masking allows for an increase in the watermark amplitude in image regions where there is significant image content, e.g. edge regions or texture regions, for the purpose of increasing the watermark robustness. The image content masks the increased amplitude so that the watermark is still below the threshold of perception. In comparison, the present invention does not use masking to maintain imperceptibility of the watermark. Instead, it uses the signal-dependent nature of imaging systems to adjust the watermark amplitude, so as to provide a visible watermark signal that is more natural in appearance. Of course, increasing the watermark amplitude in certain areas because of signal-dependent noise may also take advantage of masking to some extent, thus lowering the visibility of the signal-dependent watermark. However, the present invention is not motivated by the desire to keep the watermark below the threshold of perception. To the contrary, the present invention specifically allows watermark patterns to be perceptible, but reduces the objectionability of the visible patterns by making the watermark have visual characteristics that are consistent with the characteristics of the imaging system. As a side benefit, the signal-dependent nature of the watermark may also improve the watermark robustness by allowing overall larger amplitudes, depending on the distribution of image values (which correspondingly determine the signal-dependent scaling factor).

[0025] The second distinction between the prior art and the present invention is that the adjustment of the watermark amplitude is based upon different image attributes. In the prior art masking techniques, the amplitude adjustment is performed over a local region (e.g., an 8 x 8 block of pixels), and the variance of the signal over the block is used to control watermark amplitude. In comparison, the amplitude adjustment in the present invention is performed separately for each individual pixel in the image, and the pixel value (not a local variance) is used to control the scaling factor.

[0026] The adjustment of the watermark amplitude at each pixel location affects the estimated message image that is obtained during the extraction process. We can rewrite Eq. 2 to include the signal-dependent scaling factor $\alpha_i(x, y)$:

$$\begin{aligned}
M'(x, y) &= \left[\frac{1}{N} \sum_{i=1}^N I'_i(x, y) \right] \otimes C(x, y) \\
&= \left[\frac{1}{N} \sum_{i=1}^N (\alpha_i(x, y) [M(x, y) * C(x, y)] + I_i(x, y)) \right] \otimes C(x, y) \\
&= \left[\frac{1}{N} \sum_{i=1}^N \alpha_i(x, y) \right] [M(x, y) * C(x, y)] \otimes C(x, y) + \left[\frac{1}{N} \sum_{i=1}^N I_i(x, y) \right] \otimes C(x, y)
\end{aligned} \tag{7}$$

[0027] As before, if we assume that the carrier is designed to have uniform Fourier amplitude, then the process of correlation of the carrier on the embedded image in Eq. 7, can be reduced to:

$$\begin{aligned}
M'(x, y) &= \left[\frac{1}{N} \sum_{i=1}^N \alpha_i(x, y) \right] M(x, y) + noise \\
&= \bar{\alpha}(x, y) M(x, y) + noise
\end{aligned} \tag{8}$$

where $\bar{\alpha}(x, y)$ is the average scaling factor across all tiles at each spatial location (x, y) . If a sufficiently large number of tiles are averaged, it is expected that $\bar{\alpha}(x, y)$ would be approximately the same for every spatial location (x, y) , i.e.,

$$M'(x, y) \approx \bar{\alpha} M(x, y) + noise. \tag{9}$$

[0028] This is the same relationship that was obtained for the case of a constant scaling factor, and thus the signal-dependent watermark process does not significantly affect the extraction process if a large number of tiles are available.

[0029] Even if only a few tiles are available, the resulting $\bar{\alpha}(x, y)$ will have mainly low frequency variations. In the method by Honsinger et al., the determination of the message data from the estimated message image $M'(x, y)$ is done on a local basis, and the low frequency variations in $M'(x, y)$ due to variations in $\bar{\alpha}(x, y)$ will not greatly affect the extraction process. However, in some cases, it may be advantageous to undo the variations in $M'(x, y)$ that have been introduced by $\alpha(x, y)$. This can never be done perfectly since we only have access to the watermarked image $I'(x, y)$, not the original image $I(x, y)$, during the extraction process. Still, $I'(x, y)$ provides an estimate of $I(x, y)$, and we can use the signal dependent function $f(I'(x, y))$ to estimate $\alpha_i(x, y)$ for each pixel in each tile. Since the watermark signal is similar to a noise field, it is also possible to filter $I'(x, y)$ to provide a better estimate of $I(x, y)$. Any number of well-known noise filtering methods can be used for this purpose, e.g., lowpass linear smoothing filters or nonlinear filters such as median filters. Regardless of the method that is used to estimate $I(x, y)$, the resulting estimate can be evaluated with the signal dependent function to produce an estimate of the watermark scaling factor at each spatial location in each tile, i.e.,

$$\alpha'_i(x, y) = f(I'_i(x, y)). \tag{10}$$

This estimate can then be used to normalize each pixel of each tile of the watermarked image prior to the averaging of the tiles:

$$M'(x, y) = \left[\frac{1}{N} \sum_{i=1}^N \frac{a_i(x, y)}{a_i(x, y)} \right] M(x, y) + \text{noise} \quad (11)$$

$$\approx M(x, y) + \text{noise}$$

[0030] In addition to applying signal dependencies to the carrier amplitude, it may also be desirable to shape the carrier frequency spectrum (specifically, the Fourier amplitude spectrum) to make the embedded signal appear more like film grain. In U. S. Patent 6,044,156, referenced above, the preferred carrier possesses a flat frequency spectrum except near the DC frequency. Referring to Fig. 5, a flat frequency amplitude spectrum 46 is preferred to ensure that the message image is reconstructed with the highest fidelity. However, as described in the Honsinger patent, the carrier spectrum is shaped to include a ramp 48 in Fourier amplitude from DC frequency in order to minimize the impact of the cross-correlation of the carrier with the image data.

[0031] In contrast to the flat spectrum recommended by Honsinger et al., the shape of the Fourier amplitude spectrum of film grain noise (as well as many other imaging system noise sources) generally is a decreasing function with frequency. Fig. 6 depicts an example of a typical Fourier amplitude spectrum 50 for film grain. U.S. Patent 5,641,596, referenced above, describes methods for characterizing the spatial correlation coefficients of film grain, which is entirely equivalent to characterizing the Fourier amplitude spectrum. As shown in Fig. 7, according to one aspect of the present invention, the carrier spectrum maintains the ramp 48 in Fourier amplitude, but the remaining carrier spectrum 50' is shaped to the same form as that of film grain. Because the carrier amplitude spectrum decreases with frequency in the region that was preferably flat according to Honsinger, the impact of this shaping is to reduce the information carrying capacity of the embedding algorithm. However, given that the film grain noise can be a preferred visual attribute, increasing the amplitude (or, equivalently, increasing the visibility) of the watermark pattern can compensate for the loss of information carrying capacity. A drawback of this embodiment is the inclusion of the ramp region in the carrier spectrum does not completely simulate the appearance of film grain noise. A solution to this problem is to add an additional noise component (independent of the watermarking process) to the dispersed message image spectrum to fill in the low frequency ramp region 52 as shown in Fig. 8. This additional noise component is not an integral part of the watermarking process, but is included only to provide better visual qualities to the watermark pattern.

[0032] In a preferred embodiment of the invention, the carrier image spectrum is shaped identically to the film grain noise. Fig. 6 depicts the amplitude spectrum of this carrier. This eliminates the need to fill in the ramp region with non-information carrying noise data. However, for the reasons discussed above, this process may affect extraction performance, because correlation of the embedded image with such a carrier will produce a strong cross-correlation of the carrier with the original image content. This problem may be eliminated by defining two carrier images, an embedding carrier image and an extracting carrier image. The embedding carrier image is designed using the full frequency spectrum shape of film grain 50 as in Fig. 6, but the extracting carrier image is filtered to include the ramp 48 at low frequencies as in Fig. 7. As a result, the embedding carrier image contains the extraction carrier image in addition to the extra low frequency energy that is needed to model the film grain noise. Practically, this causes no performance problems and allows a simpler embedding implementation.

[0033] For completeness, we note that the phase spectrum of film grain noise is of random character. Visually, observers do not perceive or care about the local pixel arrangements due to the phase of the noise. They are far more sensitive to the effects produced by the amplitude spectrum of the noise. Therefore, the phase generation process of the carrier requires no change from that described in U. S. Patent 6,044,156.

[0034] As mentioned previously, the present invention offers a method for minimizing the objectionability of a high amplitude watermark. However, in some systems, suprathreshold watermarking may actually improve perceived quality. An example of such a system is digital cinema. Digital cinema refers to the electronic distribution and display of theatrical movies. Studios and distributors have a strong need to protect the movie content from unauthorized use, and watermarking can assist by establishing ownership and tracing the source of stolen content (through the use of hidden date/time/location stamps inserted at the time of the movie presentation). To facilitate the distribution of the movie content, it is typically compressed by a significant amount using a technique such as MPEG. The process of compression often removes a large portion of the high frequency content, resulting in a displayed movie without the expected look that includes film grain noise. The present invention can be used to add signal-dependent noise to the compressed movie prior to display, thus providing a more natural look, while also providing the benefit of a watermark signal.

[0035] In the preceding discussion, the present invention was used to replace film grain noise that had been removed by a compression process, while simultaneously providing a watermark signal. The same concept can be extended to provide benefits in other systems. For example, it may be advantageous to deliberately remove the inherent noise from

an image or image sequence (using one of any well-known noise reduction techniques), and then replace the removed noise with a watermark signal that has a desired visual appearance. The watermark signal may have a visual appearance that is similar to the noise that was removed, or it could have completely different appearance (for example, a higher or lower noise level, or noise characteristics that mimic a different type of film). Techniques for creating noise fields with the desired characteristics are described in the aforementioned patent by Gray et al. The removal of existing noise and subsequent replacement by a watermark signal also provides for a system where an existing watermark pattern (which typically has the appearance of a noise field) can be effectively removed and then a new watermark pattern (perhaps containing different information than the original watermark or a combination of information from the original watermark and new information) is inserted. In a conventional system, a new watermark would be added directly to the image containing the original watermark, thus elevating the overall noise level due the addition of two watermark patterns. By effectively removing the first watermark pattern prior to the insertion of the new watermark pattern, the overall noise of the twice-watermarked image is kept to a desired level.

[0036] While the invention has been discussed in terms of the watermarking process as described by Honsinger et al., it is obvious how the same method can be applied to any spatial domain watermarking process that uses a random noise carrier in the watermarking process. Some spatial domain watermarking processes use more than one carrier within a single image, where the resulting watermark pattern is the sum of the carriers after modulation by the message information. Signal-dependent scaling of the watermark pattern can also be performed for these types of watermarking methods. The invention can also be used for some types of frequency domain watermarking methods that use a PN sequence to produce a watermark pattern that has the visual characteristics of random noise. However, most frequency domain methods use frequency decompositions over local regions (e.g., 8 x 8 blocks, and it is not possible to adjust the watermark amplitude at each pixel location. Still, signal-dependent adjustment of the watermark can be performed using the average signal value within the local region, and this process is fundamentally different than the prior art methods that consider masking in adjusting the watermark amplitude.

Claims

1. A method for embedding a watermark signal that contains message data in a digital image represented as a two-dimensional array of pixel values, comprising the steps of:
 - a) providing a dispersed message image having pixel values representative of the message data;
 - b) modifying each pixel value of the dispersed message image as a function of the corresponding pixel value in the digital image; and
 - c) combining the modified dispersed message image with the digital image to produce a watermarked image.
2. The method claimed in claim 1, wherein the step of providing a dispersed message image, comprises the steps of:
 - a1) producing a message image representing the message data;
 - a2) providing a carrier image; and
 - a3) convolving the message image with the carrier image to produce the dispersed message image.
3. The method claimed in claim 2, wherein the carrier image has random phase.
4. The method claimed in claim 3, wherein the carrier image has a Fourier amplitude that models the Fourier amplitude of a specified system noise.
5. The method claimed in claim 3, wherein the carrier image has a Fourier amplitude that models the high frequencies of the Fourier amplitude of a specified system noise and contains a ramp from zero for low frequencies.
6. The method claimed in claim 1, wherein the step of modifying each pixel of the dispersed message image comprises multiplying the pixel by a scaling factor representative of a specified system signal-dependent noise.
7. The method claimed in claim 1, further comprising the step of extracting the message data from the watermarked image.
8. The method claimed in claim 2, wherein the step of modifying each pixel of the dispersed message image comprises multiplying the pixel by a scaling factor representative of a specified system signal-dependent noise.

EP 1 215 624 A2

9. The method claimed in claim 8, further comprising the step of extracting the message image from the watermarked image by correlating the carrier image with the watermarked image.
- 5 10. The method claimed in claim 9, further comprising the steps of forming an estimate of the scaling factor for each pixel value in the dispersed message image, and dividing each pixel value of the watermarked image by the corresponding estimated scaling factor prior to extracting the message data from the watermarked image.

10

15

20

25

30

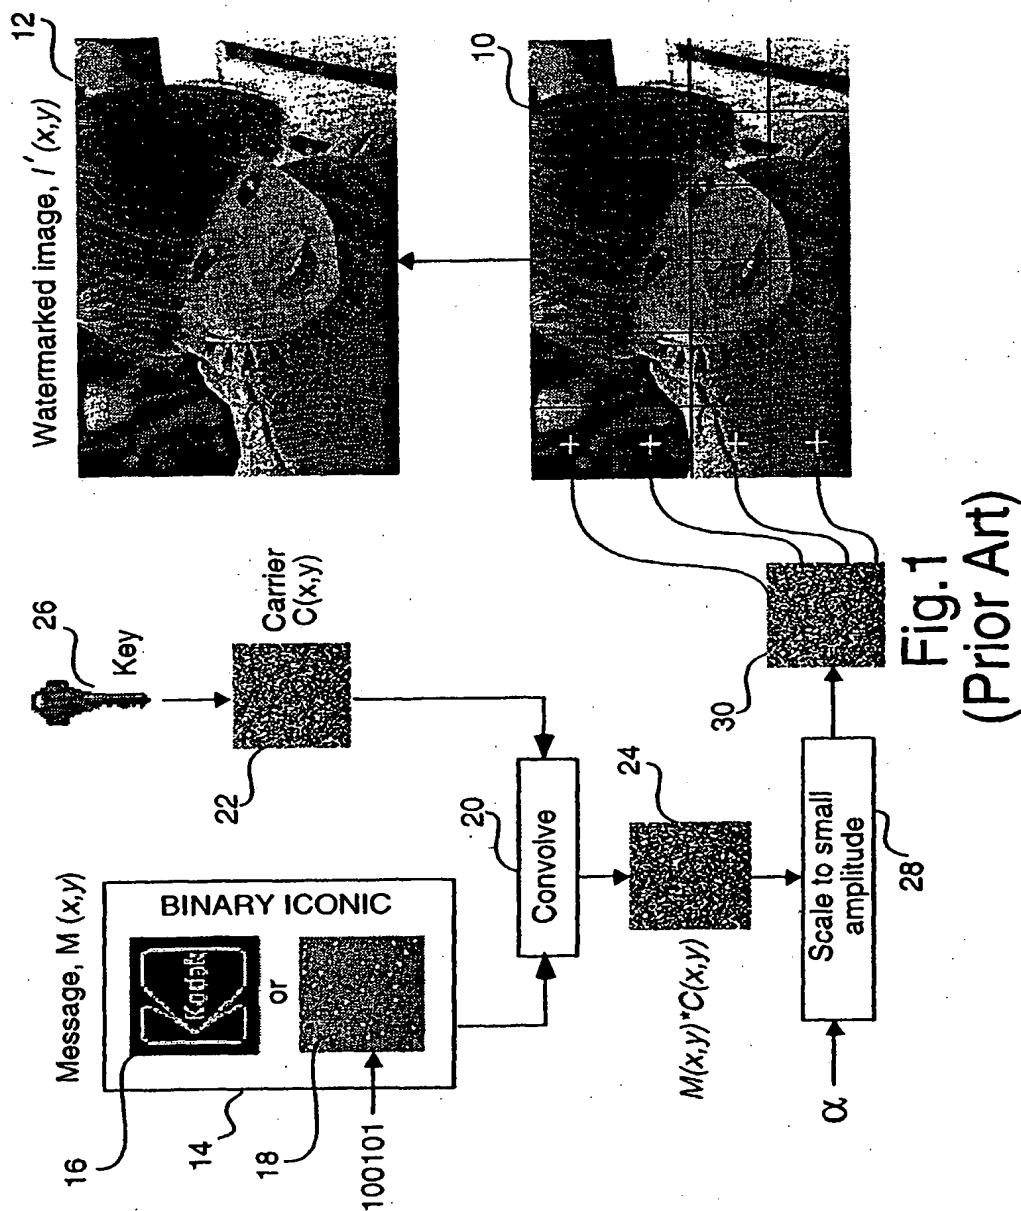
35

40

45

50

55







Fig.2
(Prior Art)






FIG. 4






FIG. 5




FREQUENCY

FIG. 6



FREQUENCY

FIG. 7



FREQUENCY

FIG. 8